

APPLICABILITY OF THE TEMPORALLY-ORDERED ROUTING ALGORITHM FOR USE IN MOBILE TACTICAL NETWORKS

Vincent D. Park

Joseph P. Macker

Information Technology Division

Naval Research Laboratory

Washington, DC

M. Scott Corson

Institute for Systems Research

University of Maryland

College Park, MD

ABSTRACT

In this paper we present a conceptual overview of the Temporally-Ordered Routing Algorithm (TORA), discuss the philosophy that shaped its design and consider its applicability for use in forward-deployed mobile tactical networks. The salient characteristics of mobile, multihop, wireless networks differ significantly from those of traditional hardwired networks. Consequently, the routing protocols that have been designed for operation in the Internet are not particularly well-suited for use in mobile tactical environments. TORA, which has been tailored for operation in this highly-dynamic networking environment, represents a significant departure from the traditional "shortest-path" routing paradigm. We also highlight recent simulation results of a performance comparison with Ideal Link-State (ILS) routing. The results show that the relative performance of TORA and ILS is critically dependent on the network size and average rate of topological changes. The results further indicate that the performance of TORA exceeds that of ILS for the conditions expected in relatively large mobile networks, lending credence to the philosophy behind the TORA design.

INTRODUCTION

Internet Protocol (IP) networking technology is primarily based on a hardwired infrastructure. Since the interconnections between the routers in a conventional IP network are hardwired, the physical topology of the network is relatively static. Thus, traditional IP routing protocols have been designed for operation in a quasi-static networking environment with hardwired links. These routing protocols are typically based on shortest-path algorithms and seek to provide least-cost paths with respect to a particular cost

metric [BG92]. This shortest-path routing paradigm is a good fit for the conventional networking environment in which it is has evolved.

A mobile, multihop, wireless network—or Mobile Ad hoc NETwork (MANET)—is far from a conventional networking environment. A MANET can be envisioned as a collection of routers (equipped with wireless receiver/transmitters) which are free to move about arbitrarily. The status of the communication links between the routers, at any given time, is a function of their positions, transmission power levels, antenna patterns, cochannel interference levels, etc. The mobility of the routers and the variability of other connectivity factors result in a network with a potentially rapid and unpredictably changing topology. In addition, wireless links inherently have significantly lower capacity than their hardwired counterparts and are more prone to congestion. Due to the considerable differences in the networking environment, the suitability of the shortest-path routing paradigm should be contemplated.

THE TEMPORALLY-ORDERED ROUTING ALGORITHM

The Temporally-Ordered Routing Algorithm (TORA) is a highly adaptive distributed routing algorithm, which has been tailored for operation in a mobile networking environment [PC97]. The basic, underlying, routing mechanism of TORA is neither a distance-vector nor a link-state algorithm; it is one of a family of "link-reversal" algorithms. A key concept in the protocol design is that it largely decouples the generation of far-reaching control message propagation from the dynamics of the network topology. This behavior makes it highly adaptive and well-suited for a dynamic mobile network with limited bandwidth. In this section, we first present the design philosophy that shaped the development of TORA, followed by a conceptual description of the protocol and some highlights of recent simulation results [PC98].

This work was supported in part by the Office of Naval Research the U.S. Army Research Laboratory's Federated Laboratory Program, Contract No. 01433109.

Design Philosophy

The design and development of a routing protocol that is better suited for operation in a MANET is a considerable challenge. An early step towards achieving the goal is to postulate—given the expected characteristics of the mobile networking environment—what properties or attributes may be desirable for a well-suited routing protocol. This is a difficult task, and the research community does not yet seem to agree on a common set of desirable properties nor on their relative importance.

An engineering approach to the problem is to design a routing protocol based on a set of postulated desirable properties and then to evaluate the relative performance of the protocol in the context of a mobile networking environment. The design choices will essentially be validated, if the newly developed protocol can be shown to outperform traditional routing approaches under the expected conditions of a mobile networking environment. In essence, this is the approach used in the development of TORA.

The following conjectures are based on the aforementioned expected characteristics of a mobile networking environment.

- Due to the potentially high rate of topological change, the protocol should *converge quickly following any reactions to topological change events*.
- Due to the limited capacity of wireless communication links (and the possible presence of energy-constrained nodes), the protocol should result in *bandwidth-efficient routing*, where bandwidth efficiency refers to *minimizing the aggregate amount of control and data traffic*.

It can easily be argued that the first attribute is important for any routing protocol, regardless of the networking environment. However, the second attribute applies principally to MANETs, where the primary system constraints are bandwidth (and possibly energy)—rather than processing and storage capacity, which limit traditional hardwired networks. TORA’s design is aimed at minimizing aggregate bandwidth demand in large, highly-dynamic wireless networks, based largely on the notion that a shortest-path routing computation may be too heavyweight for efficient operation in these systems. The idea here is that there is some *minimum communication overhead* associated with performing a shortest-path computation, and that this overhead may consume too much of the network’s bandwidth—leaving too little for data communication. A routing algorithm that forgoes such a computation in favor of a lighter-weight computation can result in less aggregate communication demand (including both control overhead and data traffic).

This concept of the “weight” of a routing computation (measured in terms of communication complexity) can be loosely described by considering the “scope” of control messaging following a topological change. An illustration of how the scope of a failure reaction (i.e., the “set” or “number” of nodes that must participate in the reaction to a link failure) can differ for various classes of routing

algorithms is depicted in Figure 1. In a link-state routing algorithm, each router maintains a complete view of network topology [BG92]. Thus, following a link failure, all nodes must be made aware of the change in link status and essentially participate in the failure reaction.




Figure 1. Scope of failure reactions.

In a path-finding algorithm, each router maintains the shortest-path spanning trees from itself and each of its neighbors to all possible destinations [CRKG89, RF89, Humblet91]. Each spanning tree consists of the distance and “predecessor” (i.e., second-to-last hop) along the shortest path from the root node of the spanning tree to each possible destination. Thus, following a link failure, the set of nodes for which the distance or predecessor to any given destination was affected by the change in link status must participate in the failure reaction. Clearly, this is less than or equal to the set of *all* nodes, which would participate in the reaction if running a link-state algorithm.

In a distance-vector algorithm, each router maintains the distances from itself and each of its neighbors to all possible destinations [BG92]. Thus, following a link failure, the set of nodes for which the distance to any given destination was affected by the change in link status must participate in the failure reaction. Again, this must be less than or equal to the set of nodes that would participate in the reaction if running a path-finding algorithm¹. It is possible for the predecessor from some node to a given destination to be affected by a link status change, while the distance to the given destination is not—in which case a path-finding algorithm would react but a distance-vector algorithm would not.

Since TORA is not a shortest-path algorithm and the state maintained by each router is significantly different, it is difficult to directly compare the scope of TORA failure reactions to those other routing approaches. In TORA, rather than maintaining “multihop topology information” or an “additive distance metric” to each destination, each router simply tries to maintain information regarding the “direction” (or set of next-hop neighbors) for forwarding packets to a given destination. Thus, a node with a “route” to a given destination has one or more of its next-hop neighbors marked

¹ This assumes that a tie-breaker is used to ensure that there is a unique shortest-path solution.

as “downstream”—where downstream paths lead to the destination. Following a link failure, only the set of nodes for which the last available downstream path to any given destination was lost due to the change in link status must participate in the failure reaction. In essence, TORA builds a multipath routing structure and uses the availability of alternate paths to limit the reactions to topological changes. Thus, it is logical that the failure reactions for TORA may be less frequent and have a smaller scope than for a distance-vector algorithm on average.

Like a distance vector algorithm, TORA maintains state on a per destination basis. This property is exploited in the design of TORA by only creating and maintaining routes *on demand*, since it may not be desirable to maintain routing between all possible source/destination pairs at all times. The overhead expended to establish a route between a given source/destination pair will be wasted if the source does not require the route prior to its invalidation due to topological changes.

Conceptual Overview

Conceptually, a logically separate version of TORA is run for each destination to which routing is required. For the following presentation, we will focus on a single version running for a given destination. TORA builds and maintains a directed acyclic graph (DAG) rooted at the destination. The DAG, by design, ensures that all directed paths are loop-free and lead to the destination. Links between routers are directed (to form the DAG) based on a metric, maintained by the routers, that can conceptually be viewed as a “height” (i.e., a link is directed from the “higher” router to the “lower” router). Given the height of a router, $H[i]$, and the height of an adjacent neighbor router, $H[j]$ —link directions are assigned as follows:

- if $H[j] == \text{NULL}$ then Unassigned
- else if $H[i] == \text{NULL}$ then Downstream
- else if $H[i] > H[j]$ then Downstream
- else if $H[i] < H[j]$ then Upstream

A conceptual illustration of the DAG for a given destination is depicted in Figure 2.




Figure 2. Conceptual illustration of the directed acyclic graph formed by the relative heights of the routers.

The protocol can be separated into three basic functions: creating routes, maintaining routes, and erasing routes. Creating a route from a given router to the destination requires establishment of a sequence of directed links leading from the router to the destination. This function is only initiated when a router with no directed links requires a route to the destination (i.e., on demand). Thus, creating routes essentially corresponds to assigning directions to links in an undirected network or portion of the network. The method used to accomplish this is an adaptation of the query/reply process described in [CE95]. Immediately following a link failure there may be directed paths that no longer lead to the destination. Maintaining routes refers to reacting to topological changes in the network in a manner such that routes to the destination are re-established within a finite time. In essence, when a router has no downstream links, it reverses the direction of one or more links by selecting a new height. The algorithm developed to accomplish this is in the same general class of algorithms presented in [GB81]. While inheriting many of the properties of the class, TORA adds the usage of “logical time” to provide an ability to detect network partitions, which leads to the third function—erasing routes. Upon detection of a network partition, all links (in the portion of the network that has become partitioned from the destination) must be undirected to erase invalid routes.

The preceding presentation of TORA is straightforward—realizing the aforementioned behavior is less so. Ensuring that the algorithm used to maintain the DAG is efficient and provides sufficient information to detect network partitions requires some subtle complexity. The height metric maintained is an ordered quintuple (τ , oid, r, δ , i) with the following values:

- τ : the “logical time” of a link failure, defining a new “reference level”
- oid: the unique ID of the router that defined the reference level
- r: a “reflection” indicator bit
- δ : a “propagation” ordering parameter
- i: the unique ID of the router

Each value in the quintuple serves a specific purpose in providing the desired functionality in the failure reaction mechanism. For a more detailed explanation of the protocol, the reader is referred to [PC97].

Performance Results

The relative performance of TORA was compared extensively with Ideal Link-State (ILS) routing and pure flooding via simulation using the Optimized Network Engineering Tool (OPNET) [PC98]. Only a cursory overview of the study and a summary of the results are presented herein. ILS was selected for comparison due to its simplicity, familiarity and the fact that ILS technology is the basis for the Open Shortest Path First (OSPF) routing algorithm [Moy94]—a widely used IP routing protocol. Flooding provided a baseline to ensure that the simulation parameter

settings and test scenarios provided a suitable networking environment.

The simulations were designed to evaluate the effect of varying the following three network characteristics:

- network size,
- rate of topological change,
- network connectivity.

A series of tests was conducted to show under what conditions TORA does and does not perform well relative to ILS, and to provide insight into its applicability for mobile wireless networks. Plots of the bandwidth utilization and end-to-end message packet delay for one of the test sequences are depicted in Figures 3 and 4.




Figure 3. Bandwidth utilization -vs.- rate of topological change.




Figure 4. Mean message packet delay -vs.- rate of topological change.

In Figure 3, the solid (lower) portion of the stacked bars represents the average number of data bits transmitted per data bit delivered (DATA), while the hashed (upper) portion represents the average number of control overhead bits transmitted per data bit delivered (CTRL). Thus, the total bar represents the *aggregate* number of bits transmitted per data bit delivered—the metric we desire to minimize. The DATA portion represents only the message packet payload bits, while the CTRL portion represents the control packet bits as well as message packet overhead (header) bits. The results

clearly indicate that as the rate of change increases the amount of control overhead increases much more rapidly for ILS than for TORA (Figure 3). In fact, at the higher rates of change depicted, *ILS utilizes more bandwidth for control overhead than for data*. As the rate of topological change increases, the increase in ILS control overhead causes an increase in the mean message packet delay (Figure 4)².

The simulations also provide insight into the effects of networking characteristics, such as network size. Overall, the simulation results indicate that the performance of TORA will eventually outperform ILS as either the:

- rate of network topological change *increases*, or
- size of network *increases*, or
- available bandwidth *decreases*.

A somewhat unexpected implication is that TORA is more scalable than ILS. However, in retrospect, this is perfectly logical due to the ability of TORA to localize algorithmic reactions.

APPLICABILITY TO TACTICAL NETWORKING

Tactical networking provides perhaps the best example of the environment in which TORA was designed to operate. Although a fixed infrastructure with hardwired links may form some parts of a networking architecture, significant portions of the network/internetwork will likely comprise mobile platforms and rely on wireless communications. This is especially true of far-forward tactical networking, where there will likely be little or no fixed infrastructure (Figure 5).




Figure 5. Visualization of a tactical network.

The architectural design of a tactical network/internetwork will largely determine where TORA is most applicable. In an architecture, where many of the mobile platforms have internal Local Area Networks (LANs) and are interconnected via wireless communications, TORA may be applicable as an internetwork routing solution. Essentially, TORA could provide equivalent functionality to that of OSPF in a typical IP internetwork. If many of the mobile routers have multiple wireless interfaces based on different wireless communications technology, this solution can also serve to

² The mean packet delay has been plotted on a logarithmic scale to provide visual separation between the data points when the delay was less than 10 seconds.

bind together the heterogeneous wireless communications infrastructure. An IP-compliant version of TORA is under development and is being promoted for potential standardization through the MANET working group of the Internet Engineering Task Force (IETF).

There is also the potential to apply TORA to an individual multihop wireless LAN that is based on a single wireless communications technology. Such a wireless LAN may or may not be part of a larger internetwork. In this case TORA could essentially be applied at the Media Access Control (MAC) level, to provide a multihop forwarding capability within the LAN.

In either of these cases, TORA may allow for larger routing domains due to the increased scalability and adaptivity of the protocol. The potential benefits of this are two-fold. First, this tends to relax the range of mobility restrictions on the mobile platform. The ability for a mobile platform to become separated from its routing domain and establish connectivity elsewhere in the internetwork requires additional complexity and overhead. By increasing the coverage area of the mobile routing domains, the need and/or likelihood of this type of roaming support can be reduced. Secondly, using a smaller number of larger routing domains may also reduce the complexity of the networking architecture—i.e., using fewer routing domains can reduce the number and type of border gateways required. The potential to support larger routing domains also means that a greater variability in domain size may be supported—a potentially important characteristic if domains can cover several mobile tactical formations, which may need to merge and recombine in unpredictable patterns during combat.

There are currently several unresolved issues. If IP networking is used, attention will need to be given to address assignments and subnet masking. The best approach for address allocation will likely be dependent on the specific architecture. Also, since it is unlikely that TORA would be used in all portions of an internetwork, solutions for gateways between TORA and other routing domains will also be required. For example, if a mobile network is connected to a larger fixed infrastructure—it may be desirable to use TORA in the mobile network but not in the fixed infrastructure. In essence, some routers will need to serve as border gateways between the different routing domains. At some point, it will also be beneficial to develop an interface between TORA and the Border Gateway Protocol (BGP), which is often used for routing between Autonomous Systems (AS) in the Internet.

CONCLUSIONS

TORA is a highly adaptive distributed routing algorithm, which has been tailored for operation in a mobile networking environment. In its design, the ability to perform a shortest-path routing computation is sacrificed for a greater ability to limit the scope of control messaging (following link failures and additions) to a small set of nodes near the change. This behavior makes it scalable, adaptive and well-suited for a dynamic mobile network with limited bandwidth. The

validity of this design choice is supported by simulation results, which indicate that the performance of TORA exceeds that of ILS—a well known and proven technology for supporting shortest-path routing—for the conditions expected in a relatively large mobile networks.

TORA is directly applicable to the tactical networking environment, and can provide either an internetwork routing solution or a multihop wireless LAN routing solution. In fact, tactical networking provides perhaps the best example of the environment in which TORA was designed to operate. Due to the scalability of the protocol, TORA may be able to support larger routing domains within the networking architecture—thus, relaxing the range of mobility restrictions on the mobile platforms within a given routing domain. Furthermore, using a smaller number of larger routing domains may also serve to reduce the complexity of the networking architecture. Overall, TORA has the potential to play an important role in the tactical networking environment.

REFERENCES

- [BG92] D. Bertsekas and R. Gallager, *Data Networks*, (Prentice-Hall, 1992).
- [CE95] M.S. Corson and A. Ephremides, “A Distributed Routing Algorithm for Mobile Wireless Networks,” *Wireless Networks* 1, (1995).
- [CRKG89] C. Cheng, R. Riley, S. Kumar, and J.J. Garcia-Luna-Aceves, “A Loop-Free Extended Bellman-Ford Routing Protocol without Bouncing Effect,” *Proc. ACM SIGCOMM '89*, Austin, TX, September 1989.
- [GB81] E. Gafni and D. Bertsekas, “Distributed Algorithms for Generating Loop-free Routes in Networks with Frequently Changing Topology,” *IEEE Trans. Commun.*, January 1981.
- [Humble91] P. Humblet, “Another Adaptive Distributed Shortest Path Algorithm,” *IEEE Trans. Commun.*, June 1991.
- [PC97] V. Park and M.S. Corson, “A Highly Distributed Routing Algorithm for Mobile Wireless Networks,” *Proc. of IEEE INFOCOM '97*, Kobe, Japan, April 1997.
- [PC98] V. Park and M.S. Corson, “A Performance Comparison of the Temporally-Ordered Routing Algorithm and Ideal Link-State Routing,” *Proc. of IEEE Symposium on Computers and Communications '98*, Athens, Greece, June 1998.
- [Moy94] J. Moy, “*OSPF Version 2*,” Internet RFC 1583, March 1994.
- [RF89] B. Rajagopalan and M. Faiman, “A New Responsive Distributed Shortest-Path Routing Algorithm,” *Proc. ACM SIGCOMM '89*, Austin, TX, September 1989.